




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶: G06F 17/30	A2	(11) International Publication Number: WO 98/58330 (43) International Publication Date: 23 December 1998 (23.12.98)
(21) International Application Number: PCT/SE98/01119 (22) International Filing Date: 10 June 1998 (10.06.98) (30) Priority Data: 08/876,588 16 June 1997 (16.06.97) US (71) Applicant: TELEFONAKTIEBOLAGET LM ERICSSON [SE/SE]; S-126 25 Stockholm (SE). (72) Inventors: MATTSOON, Håkan; Brittsommargränd 2, S-135 61 Tyresö (SE), WIKSTRÖM, Claes; Sättagårdsvägen 73, S-127 36 Skårholmen (SE). (74) Agent: ERICSSON TELECOM AB; IPR Management & Patent Dept., S-126 25 Stockholm (SE).		(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, GW, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TO). Published <i>Without international search report and to be republished upon receipt of that report.</i>

(54) Title: LINGERING LOCKS FOR REPLICATED DATA OBJECTS**(57) Abstract**

In a network (20) wherein replicas of a data object reside at a plurality of nodes (30A, 30B), a node has a lock manager (73) which handles requests for access to the data object by processes of differing nodes. When a process native to the node desires access to the data object, the lock manager sets a lingering lock flag in a lock table (220) for the data object. Setting the lingering lock flag indicates that the data object can only be accessed at the node to the exclusion of processes of other nodes. The lingering lock flag can remain set after the process native to the node terminates access to the data object, thereby enabling a successive access to the data object by the same process or another process native to the node, such successive access being without notification to other nodes of the network.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon		Republic of Korea	PT	Portugal		
CN	China	KR	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

LINGERING LOCKS FOR REPLICATED DATA OBJECTS

BACKGROUND

This invention is related to that described in simultaneously-filed U.S. Patent Application Serial No. 08/____,____, (attorney docket 1410-316), entitled "INITIALIZATION OF REPLICATED DATA OBJECTS", which is
5 incorporated herein by reference.

1. Field of Invention

This invention pertains to distributed or replicated databases, and particularly to the locking of
10 access to data objects maintained by such databases.

2. Related Art and Other Considerations

When executing application programs, computers frequently assign and/or change values of data objects. In
15 many instances the data objects are stored as part of a database. In more complex systems in which a plurality of computers are networked together, more than one computer may require access to a certain data object, and may change or update the value of that data object. To cater to such a
20 multi-computer networked environment, a replicated database system can be established. In a replicated database system, each computer can maintain its own version of the database so long as all other computers are advised of the changes of a data object for maintaining consistency among the
25 replicated databases.

Replicated databases have two primary advantages. A first advantage is fault tolerance. A second advantage is that local access to a replicated database is faster and less expensive than remote access to a database at another computer.

Each computer can have a plurality of application programs or processes accessing its local version of a replicated database. When a particular process needs to access a local version of the replicated database, a "lock" of the data object must be obtained by the accessing process in order to ensure exclusive access to the data object. Many so-called "concurrency control" algorithms or strategies use various forms of locking to achieve this goal.

The problem with locking of data objects in a network context is the increased effort and traffic occasioned by generating and handling lock-related messages between various nodes of the network. Consider, for example, a multi-node network which includes nodes N1 and N2, each having a computer which executes many processes and a local version of a replicated database in which resides data object Y. The computer of each node has a special purpose process known as a lock manager for administering the setting and releasing of locks for data objects including object Y. When a first process of node N1 desires access to data object Y, the computer of node N1, on behalf of the first process of node N1, sends a lock request message to the lock managers of all nodes of the network having data object Y in their versions of the replicated database. In response, each node in the network returns a lock status message to node N1, e.g., confirming (when

appropriate) that data object Y is now be locked at node N1. The first process of node N1 then can access and (when necessary) modify the value for data object Y. When the first process is finished with data object Y, the computer
5 of node N1 on behalf of its first process sends a lock release message to all nodes. If a second process at node N1 then desires access to data object Y, essentially the same steps as described are executed, including issuance of the lock request message, collecting of response messages,
10 and (upon completion of usage of data object Y) issuance of lock release messages.

Thus, repeated requests by processes at the same node for locking a data object can entail considerable
15 message traffic over the network, as well as complicated lock bookkeeping by lock managers situated throughout the network. Various strategies to counter these undesirable ramifications have been developed, as explained, for example, by Bernstein, P.A., et al., *Concurrency Control and*
20 *Recovery In Database Sytems*, Addison Wesley, 1987; and Özsu, T., Valduriez, P., *Principles of Distributed Database Systems*, Prentice Hall, 1991.

One strategy is merely to ignore the overhead
25 associated with network input/output when setting locks on remote objects. A second strategy is to organize the data as having one primary copy and a set of secondary copies. With such organization, the secondary copies are not visible to the programs or processes accessing the data. The
30 secondary copies are updated by the system whenever the primary copy is updated. The major disadvantage of this second strategy is that the system becomes static. All programs are forced to use the primary copy and, if the

primary copy fails, or if for any reason one of the secondary copies wants to act as the primary copy, the system must be reconfigured.

5 A third strategy is to employ a loose or non-existent locking in which multiple copies can be updated simultaneously without prior locking. This third strategy works primarily in certain peculiar situation, e.g., when different operations on the same data object commute. For 10 example, if an integer value is stored, and the only operation allowed on the data object is incrementation, i.e., a counter, locking is not necessary.

What is needed therefore, and an object of the
15 present invention, is simplified method and apparatus for
locking access to data objects of a database which is
replicated across a network.

20 SUMMARY

In a network wherein replicas of a data object reside at a plurality of nodes, a node has a lock manager which handles requests for access to the data object by processes of differing nodes. When a process native to the node desires access to the data object, the lock manager sets a lingering lock flag in a lock table for the data object. Setting the lingering lock flag indicates that the data object can only be accessed at the node to the exclusion of processes of other nodes. The lingering lock flag can remain set after the process native to the node terminates access to the data object, thereby enabling a successive access to the data object by the same process or

another process native to the node, such successive access being without notification to other nodes of the network.

The lingering lock flag is in addition to an
5 object lock flag. When the process native to the node desires access to the data object, the lock manager also sets the object lock flag in the lock table for the data object. When the process native to the node terminates
10 access to the data object, the lock manager clears the object lock flag.

The lock manager also handles a lock scheduling queue for the data object. When the lock manager receives a request from a non-native process from a second node
15 relating to a desired access of the data object at the second node while the lingering lock flag is set for the data object in the lock table of the node, the lock manager can list the non-native process of the second node as requesting access to the data object in the lock scheduling
20 queue.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, features, and advantages of the invention will be apparent from the
25 following more particular description of preferred embodiments as illustrated in the accompanying drawings in which reference characters refer to the same parts throughout the various views. The drawings are not necessarily to scale, emphasis instead being placed upon
30 illustrating the principles of the invention.

Fig. 1 is a schematic view of a network including two nodes whereat replicated data objects reside.

Fig. 2 is a schematic view of a node of the network of Fig. 1.

Fig. 3A is a schematic view of a series of data object accesses according to the present invention in a first scenario.

Fig. 3B is a schematic view of a series of data object accesses according to the present invention in a second scenario.

Fig. 4 is a flowchart showing steps executed by a lock manager process upon receipt of a lock-related message from another process.

Fig. 5 is a state diagram showing states of a process which sends lock-related message to lock manager process(es).

DETAILED DESCRIPTION OF THE DRAWINGS

In the following description, for purposes of explanation and not limitation, specific details are set forth such as particular architectures, interfaces, techniques, etc. in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced in other embodiments that depart from these specific details. In other instances, detailed descriptions of well known devices, circuits, and methods are omitted so as not to obscure the description of the present invention with unnecessary detail.

Fig. 1 shows a network 20 comprising two illustrative nodes 30A and 30B, also referenced separately as nodes "A" and "B" and generically as "node 30". Each node 30 has its own version of a data object X. Specifically,
5 node 30A has hard disk 32A whereon its version of data object X, referenced as X-A, is stored. Similarly, node 30B has hard disk 32B whereon its version of data object X, referenced as X-B, is stored. As employed herein, "data object" can refer to a single data value, or to a collection
10 or table of data values such as often occurs in a database. Each object (or group of objects) has its own object identifier, Oid.

Since versions of data object X are stored both at
15 nodes 30A and 30B, when one node updates the value of data object X, the updated value is communicated to the other node so that the other node can likewise have the updated value, thereby maintaining a coordination of the value of the data object X.

20 Each node 30 includes a processor or CPU 40 which is connected by an internal bus 42 to numerous elements. Illustrated ones of the elements connected to internal bus 42 include a read only memory (ROM) 43; a random access
25 memory (RAM) 44; a disk drive interface 45; and, a network interface 46. Disk drive interface 45 is connected to disk drive 50. Network interface 46 is connected to network link 60 over which the nodes 30A and 30B communicate.

30 Hard disk 32 is one example of a node-status inviolable memory or storage medium. By "node-status inviolable" is meant that the contents of the memory remain unaffected when the node crashes or assumes a "down" status.

Although the node-status inviolable memory is illustrated in one embodiment as being a hard magnetic disk, it should be understood that other types of memory, e.g., optical disk, magnetic tape, etc., are also included.

5

Processor 40 executes a set of instructions in an operating system, which in turn allow processor 40 to execute various application programs or processes 70 which are preferably stored on hard disk 32. Of particular
10 interest to the present invention is a set of instructions embodied in a computer product and known as a lock manager application program (LOCK MANAGER) 73. The effect of operation of lock manager application program (LOCK MANAGER) 73 is hereinafter described.

15

Fig. 2 shows node 30A in more detail and from a functional perspective. It will be understood that comparable structure and functionality also exists at other nodes of network 30, e.g., node 30B.

20

Processor 40A of node 30A is now described as being illustrative of processors of nodes throughout network 20. Processor 40A executes both application programs or processes 70 and the lock manager application program (LOCK
25 MANAGER) 73. In order to be executed, such processes and programs must be loaded from the medium on which they are stored, e.g., hard disk 32, into RAM 44A, for which reason exemplary processes Pid A-1 and Pid A-2 are shown as enveloped by broken line 44A in Fig. 2.

30

Processor 40A of node 30A, when executing the lock manager application program (LOCK MANAGER) 73, performs the various functions enveloped by broken line 73 in Fig. 2.

Such functions performed by the lock manager application program (LOCK MANAGER) 73 include those of message encoder/decoder 100, message interpreter 200, message generator 202, data object transaction monitor 204, lock table supervisor 206, lock schedule queue handler 208, and confirmation coordinator 210.

In addition to having instructions for the lock manager application program (LOCK MANAGER) 73 and processes
10 Pid A-1 and Pid A-2, RAM 44A has stored therein data object location table 218; object lock table 220; sticky lock table 221; and, a set lock scheduler queue 222. Collectively, object lock table 220 and sticky lock table 221 are referenced herein as the "lock table".

15

As illustrated in Fig. 2, object lock table 220 has a record 230 for each data object pertinent to node A. Each record 230 has a field 232 which identifies the data object and a field 234 wherein is stored an object lock
20 flag. The object lock flag is an identification of a particular process which has requested a lock on the data object.

Sticky lock table 221, also known as the
25 lingering lock table, likewise has a record 240 for each data object entered therein. Each record 240 has a field 242 which identifies the data object and a field 244 wherein is stored a sticky lock flag, also known as the lingering lock flag. The value set for the sticky lock flag is the
30 node where the sticky lock has been applied.

-10-

As also illustrated in Fig. 2, the lock scheduler queue 222 is comprised of a linked list of records 250. Each record 250 of queue 222 includes a field 252 which identifies the data object and a field 254 wherein is stored an identification of a requesting process, i.e., a process requesting a lock. As explained hereinafter, the lock scheduler queue 222 is essentially a first-come, first-serve listing of processes in this node or other nodes which request access to the data object which is the subject of the queue.

Operation of lock manager application program (LOCK MANAGER) 73 for node 30A, hereinafter simply referred to as the "lock manager LM", is understood in connection the flowchart of Fig. 4. Lock manager application program (LOCK MANAGER) 73 is a separate process, which performs the steps shown in Fig. 4 upon receipt of a lock-related message from another process executing at node 30A. The lock-related messages generated by a process can be one of five message types: [1] an "object_lock request" message (also known as a regular "lock request" message) having the format {stick, Oid, Node}; [2] an "unlock request" message having the format {unlock, Oid}; [3] a "test_sticky_lock" message having the format {test_sticky_lock, Pid, Oid}; [4] a "sticky_lock request" message having the format {stick, Oid, Node}; and, [5] an "unstick request" message having the format {unstick, Oid}. In each of the message formats, it is understood that the first parameter thereof pertains to the message type, that "Oid" refers to the data object identifier, that "Pid" refers to an identifier of the process sending the message; and that "Node" refers to a node of the network.

-11-

If, at step 4-2, the message is determined to be an "object_lock request" message, execution branches to a "try_lock" procedure which comprises steps framed by broken line 400 in Fig. 4 (i.e., steps 4-3 through 4-6). At step 5 4-3, the lock manager LM checks object lock table 220 to determine whether the data object identified in the message is already locked. If the data object is not already locked, at step 4-4 lock table supervisor 206 updates object lock table 220 by preparing a record 230, inserting therein 10 the Oid of the request message in field 232 and setting the object lock flag by storing the Pid of the requesting process in field 234 of the record. Then, as indicated by step 4-5, the lock manager LM sends a "granted" message to the requesting process.

15

If, at step 4-3, it is determined that the message is already locked, at step 4-6 lock table supervisor 206 directs lock schedule queue handler 208 to append a record reflecting the request in lock scheduler queue 222. The 20 record 250 entered into lock scheduler queue 222 identifies the data object and the process which requested the lock.

Should the message decoded be determined at step 4-7 to be an "unlock request" message, at step 4-8 the lock 25 manager LM deletes the record for the data object identified in the "unlock request message" from the object lock table 220. Then, if there are any entries in the lock scheduler queue 222, lock table supervisor 206 working together with lock schedule queue handler 208, updates both object lock 30 table 220 and lock scheduler queue 222. In this regard, at step 4-9 the highest priority lock request record in lock scheduler queue 222 is effectively transferred into object lock table 220, thereby enabling the awaiting process in

queue 222 to place its lock. Then, at step 4-10, the lock manager LM sends a "granted" message to the process which has just been afforded the lock by the record transfer to object lock table 220.

5

Should the message decoded be determined at step 4-15 to be a "test_sticky_lock" message, at step 4-16 the lock manager LM determines whether a lock for the data object identified in the message has a sticky lock applied thereto. In particular, lock table supervisor 206 searches the sticky lock table 221 to determine if the Oid of the message has its lingering lock flag set, i.e., a sticky lock. The lingering lock flag is considered to be "set" when field 244 for the record for the Oid has a node value stored therein. If a lingering lock flag has not been set, at step 4-17 lock table supervisor 206 apprises message generator 202, and a "not_stuck" message is sent to the requesting process.

20

If it is determined at step 4-16 that a lingering lock flag has been set for the Oid of the "test_sticky_lock" message, at step 4-18 lock table supervisor 206 determines whether the lingering lock flag is set by a process at this node or a process of another node. Such determination involves checking field 234 for the Oid of the message. If the lingering lock flag is set by a process at another node, message generator 202 is instructed to send a "stuck elsewhere" message to the process which sent the test sticky lock message (see step 4-19). If the lingering lock flag is set by a process at this node, the lock manager LM performs the "try_lock procedure" 400 above-described with reference to steps 4-3 through 4-6.

30

-13-

Should the message decoded be determined at step 4-25 to be a "sticky lock request" message, at step 4-26 the lock table supervisor 102 of the lock manager LM creates a record 240 in sticky lock table 221. In creating the

5 record, lock table supervisor 220 inserts the Oid of the data object to be locked in field 242, and the node of the requesting process in field 244. Setting the node of the requesting process in field 244 amounts to setting the lingering lock flag, also known as the sticky lock flag, for

10 the data object identified in the sticky lock request message.

Should the message decoded be determined at step 4-30 to be an "unstick request" message, at step 4-31 the

15 lock table supervisor 102 of the lock manager LM deletes the record in sticky lock table 221 for the data object identified in the "unstick request" message. This amounts to clearing the lingering lock flag, also known as the sticky lock flag, for the data object identified in the

20 "unstick request" message.

Whereas Fig. 4 shows steps executed by the lock manager LM upon receipt of a lock-related message from another process, Fig. 5 shows states of a process, such as

25 process Pid A-1, which sends lock-related message to lock manager process(es) [LMs]. State 5-0 depicts normal execution of the process. Various other states of Fig. 5 involve messages including those having formats discussed above in connection with Fig. 4.

30

State 5-1 reflects the process sending a "lock request" message to all lock managers (LMs). As indicated by state 5-2, execution by the process is suspended until a

-14-

"granted" message is received from all lock managers. When the process desires to unlock a data object, an "unlock request" message is sent as depicted by state 5-3, without disrupting execution of the process.

5

State 5-4 depicts the process sending a "test_sticky_lock" message to its local lock manager. As described above with reference to Fig. 4, three possible responses may be received from the local lock manager. A first response may be a "granted" message (see step 4-5), which results when the sticky lock flag has already been set at this node. A second response is a "not_stuck" message (generated at step 4-17 of Fig. 4). A third response is a "stuck_elsewhere" message (generated at step 4-19 of Fig. 4).

15

Upon receipt of a "not_stuck" message, as indicated by state 5-5 the process sends a (object) "lock request" message to all lock managers. The "lock request" message is handled by the lock managers in accordance with the "try_lock" procedure 400 (see Fig. 4). This results in a "granted" message being sent to the process. Then, as indicated by state 5-6 the process sends a "sticky lock request" message to all lock managers. All lock managers then update their sticky lock tables 221 to insert a sticky lock flag, as previously explained with reference to step 4-26.

Upon receipt of a "stuck_elsewhere" message, as indicated by state 5-7 the process sends a "lock request" message to all lock managers. The "lock request" messages are handled by the lock managers in accordance with the "try_lock" procedure 400 (see Fig. 4). State 5-8 shows the

30

-15-

process awaiting receipt of a "granted" message from all lock managers. Upon receipt of a "granted" messages from all lock managers, the process sends either an "unstick request" message (states 5-9) to all lock managers.

5

Operation of lock manager application program is illustrated in connection with various scenarios represented by Fig. 3A and Fig. 3B. For sake of simplicity, the nodes are illustrated primarily as including only lock managers, as well as various processes executed at those nodes (e.g., processes Pid A-1 and Pid A-2 at node 30A; processes Pid B-1 and Pid B-2 at node 30B). Fig. 3A and Fig. 3B also depict in simplified form the contents of object lock table 220, sticky (lingering) lock table 221, and lock scheduler queue 222, for each of the lock managers for the respective nodes. Only the portion of the contents of the lock tables relevant to data object X are shown, with such records of the lock tables being shown essentially in a time sequenced format.

20

As understood from the detailed discussion provided below, the present invention utilizes two types of lock flags. A first type of lock flag, known as the object lock flag, is set in the field 234 of the record 230 of object lock table 220 for the respective data object as the object lock flag. A second type of lock flag, known as the lingering lock flag or "sticky lock" flag, is set in the manner hereinafter described in field 236 of the record 240 of sticky lock table 221 for the respective data object.

30

In Fig. 3A, it is assumed that initially (event E0) no node has set a lingering lock flag for data object X. Nor has any process set an object lock with respect to data object X. Fig. 3A shows depiction of contents of each of

-16-

the object lock flag, sticky lock flag, and scheduler queue 222 at various event times hereinafter noted. Initially, i.e., at event E0, no lock flag has been set and the scheduler queue is empty.

5

However, process Pid A-1 desires to set a sticky lock flag. To do so, and as indicated by event E1 of Fig. 3A, process Pid A-1 sends a "test_sticky_lock" message to its local lock manager LMA (state 5-4 of Fig. 5). The data
10 object transaction monitor 204 of the local lock manager LMA detects such message (step 4-15) and sends a command for lock table supervisor 206 to analyze its sticky lock table 221. In the present scenario, the lock table supervisor determines (step 4-16) that no sticky lock flags have been
15 set for data object X. In accordance with such determination, lock table supervisor 206 sends a command to message generator 202, the command having a command type that causes generator 202 to send a "not_stuck" message to process Pid A-1 as event E2 (see step 4-17).

20

Upon receipt of the "not_stuck" message from lock manager LMA, process Pid A-1 requests locking of data object X by sending a "lock request" message to its local lock manager (event E3, state 5-5). When data object transaction
25 monitor 204 of the local lock manager determines that process Pid A-1 has requested a lock, a command is sent to lock table supervisor 206 and to message generator 202. In response, lock table supervisor 206 sets both the object lock flag in field 234 of object lock table 220 (step 4-4).
30 In addition, message generator 202 prepares a "granted" message which is sent as event E4 to process Pid A-1 (step 4-5). Fig. 3A shows the object lock flag set at event E4.

-17-

Upon receipt of the "granted" message, process Pid A-1 acquires state 5-6 in which a "sticky lock request" message is sent to the lock managers of each of the nodes whereat data object X is replicated. Event E5 shows the

5 "sticky lock request" message being sent first to the local lock manager LMA. The local lock manager LMA accesses data object location table 218 in order to determine what other nodes of network 20 have the data object X at replicated databases. Upon consultation of data object location table

10 218, message generator 202 directs message encoder/decoder 100 to prepare a message (which is applied to network link 60 by network interface 46A) requesting that other nodes also lock the data object. In the present illustration, node B also has data object X, for which reason event E5B

15 shows a network lock message being transmitted to lock manager LMB of node B.

Upon receipt of the "sticky lock request" message, the lock managers of nodes A and B prepare a record in their

20 respective sticky lock tables 221A, 221B for data object X, and set the sticky lock flag by storing an indication of node A in field 244 of such records in accordance with step 4-26. When lock table supervisor 206 sets the object lock flag to "Pid A-1" and the lingering lock flag to "A", lock

25 table supervisor 206 sends a confirmation message to confirmation coordinator 210. Likewise, the lock manager LMB for node B sends a "sticky lock confirmation message" over the network to node A (see event 6B). At node A, the lock confirmation message from node B is decoded by message

30 encoder/decoder 100 and applied to message interpreter 200. Upon discerning that the message is a lock confirmation message, message interpreter 200 forwards the confirmation to confirmation coordinator 210. Noting confirmation from

-18-

both lock table supervisor 206 of node A and from all other nodes whereat data object X is replicated, confirmation coordinator 210 authorizes message generator 202 to send a "sticky lock confirmation" message to process Pid A-1 as event E6. Fig. 3B show that, by completion of event E6, the sticky lock flags have been set in both sticky lock table 221A of node A and sticky lock table 221B of node B.

Upon receipt of the "sticky lock confirmation" message, process Pid A-1 is qualified to access data object X, to the exclusion of all other processes both at node A and any other nodes.

When process Pid A-1 is finished with data object X, in the particular scenario illustrated in Fig. 3A the Pid A-1 enters state 5-3 and sends an "unlock request" message as event E7 to its local lock manager LMA. Upon receipt of the "unlock request" message, lock manager commands the lock table supervisor 206 to delete the object lock flag from object lock table 220 (step 4-8). The lingering lock flag remains set at the time of event E7 as shown in Fig. 3A.

Of course, upon release of data object X by process Pid A-1, other nodes of the network having a replication of data object X are sent messages enabling those other nodes to update their local versions of data object X in accordance with the value(s) assigned or modified by process Pid A-1. However, the lock tables for such other nodes remain unchanged. That is, there is no network message to other nodes of the network reflecting a change of lock status.

Fig. 3A further illustrates that, after release of data object X by process Pid A-1, process Pid A-2 desires access to data object X and enters state 5-1 to send as event E8 a "lock request" message to lock manager LMA of node A. The requested access for data object X is detected by data object transaction monitor 204, and a command sent to lock table supervisor 206. Lock table supervisor 206 determines that the object lock flag is not set (step 4-3).

Such being the case, lock table supervisor 206 sets the object lock flag in field 234 to indicate process Pid A-2, and as event E9 the lock manager LMA sends a "granted" confirmation message (step 4-5) to process Pid A-2. Fig. 3A shows, for event E8, the object lock flag of node A being set to indicate process Pid A-2.

15

Upon receipt of the "granted" message of event E9, process Pid A-2 is able exclusively to access and modify data object X, in like manner as done by process Pid A-1 previously. When process Pid A-2 releases access to data object X (see event E10), a similar procedure as above described with respect to event E7 occurs. Included in such procedure is lock table supervisor 206 clearing the object lock flag in field 234 of object lock table 220. Again, the lingering lock flag in field 244 is not cleared in the lingering lock table of any node.

25

From the foregoing it can be seen that setting of the lingering lock flag, also known as the sticky lock flag, expedited access by process Pid A-2 to data object X.

Whereas, in the prior art, further network messages would be necessary in order for process Pid A-2 to claim exclusive access to data object X and to receive confirmation of such exclusive access, such further network messages are obviated

30

in the present invention. In other words, once a lingering lock flag has been set for a node, the lingering lock flag remains stuck at that node (until removed by operations subsequently described herein). The sticking of the

5 lingering lock flag enables the data object to be accessed repeatedly by the same process at such node, or by other processes of that node, without issuance and reception of network messages such as lock request and lock confirmation network messages. The lingering lock flag and its

10 employment as described herein thus advantageously curtails network traffic in situations in which a same node has process(es) which make successive access to the same data object.

15 Fig. 3B illustrates another scenario in which process Pid A-2 of node A has requested access and received lock confirmation with respect to data object X, with the lingering lock flag having previously been set for node A in the manner described above with respect to Fig. 3A. In

20 essence, event E11 and event E12 of Fig. 3B correspond to events E8 and E9, respectively, of Fig. 3A. Accordingly, at event E12 the object lock flag is set to indicate process Pid A-2 and the lingering lock flag remains set for node A. Likewise, as understood from the foregoing discussion of

25 Fig. 3A, the lingering lock flag for node B remains set as shown to indicate that the lingering lock is at node A.

In the scenario of Fig. 3B, however, process Pid B-1 of node B enters state 5-4 and requests access to its

30 local version of data object X while process Pid A-2 is still utilizing data object A. The request is in the form of a "test_sticky_lock" message which is sent as event E13.

-21-

Upon determining that the lingering lock is set in its sticky lock table 221B to indicate node A (steps 4-16 and 4-18), the lock table supervisor 206 of lock manager LMB causes its message generator 202 to send (as event E14) a
5 "stuck_elsewhere" message (step 4-19) to process Pid B-1.

Upon receipt of the "stuck_elsewhere" message, process Pid B-1 enters state 5-7 and as event E15 sends a "lock request" message to lock managers at all nodes whereat
10 data object X resides. As explained initially, process Pid B-1 initially sends the "lock request" message to its local lock manager, and the local lock manager sends the lock request manager to other nodes as indicated by event E15A.

15 After sending the "lock request" message of state 5-7, process Pid B-1 assumes state 5-8 and awaits a "granted" message from each node whereat data object X resides in a replicated data base.

20 Upon receipt of the "lock request" message, lock manager LMB checks its object lock table 220B. Failing to find any indication of an object lock (step 4-3), lock manager LMB updates its object lock table 220B whereby its object lock flag shows process Pid B-1 as having the object
25 lock (step 4-4). Then, lock manager LMB causes its message generator 202 to send a "granted" message (step 4-5) to process Pid B-1 as event E16.

Upon receipt of the "lock request" message, lock
30 manager LMA also checks its object lock table 220A, and finds (at step 4-3) that local process Pid A-2 still has the object lock. At this point, lock manager LMA cannot relinquish the lock on data object X. Rather, lock manager

-22-

LMA instructs its lock schedule queue handler 208 to store, in its lock scheduler queue 222A, an indication that process Pid B-1 of node desires access to data object X.

5 Process Pid B-1 thus remains suspended in state 5-5, awaiting receipt of a "granted" message from all nodes. The "granted" message has thus far been received only from its own local lock manager LMB; a "granted" message is still lacking from node A.

10

 It so happens in the scenario of Fig. 3B that process Pid B-2 then also requests access to data object X. In particular, upon assuming state 5-1 the process Pid B-2 sends an object "lock request" message as event E16 to lock
15 managers at all nodes where data object X resides. As explained previously, the object "lock request" message is initially sent to local lock manager LMB (event E17), which prepares a network message for sending of the object "lock request" message to other nodes (event E17A).

20

 At this juncture process Pid A-2 is still utilizing data object X, and process Pid B-1 preceded process Pid B-2 in requesting access to data object X.

25

 Upon receipt of the object "lock request" message of event E17 for access to data object X by process Pid B-2, lock manager LMB notes from its object lock table 220B that process Pid B-1 has the object lock flag set in field 234 (step 4-3). Such being the case, at step 4-6 the lock
30 manager LMB has its lock schedule queue handler 208 store in queue 222B an indication that process Pid B-2 next desires access to data object X. Process Pid B-2 must then assume state 5-2 of awaiting a yet-issued "granted" message from

all nodes whereat data object X is replicated. A similar activity occurs at node A, whereat in accordance with steps 4-3 and 4-4 the process Pid B-2 is added to lock scheduler queue 222A.

5

- Eventually process Pid A-2 finishes with data object X. Process Pid A-2 assumes state 5-3 and sends an "unlock request" message to its local lock manager LMA, as indicated by event E18 of Fig. 3B. Upon receipt of the
- 10 "unlock request" message, lock manager LMA deletes the indication of process Pid A-2 from its object lock table 220A (step 4-8). Moreover, as indicated by step 4-9, lock manager LMA transfers the record in its lock scheduler queue 222A which refers to process Pid B-1 to its object lock
- 15 table 220A, setting the object lock flag in field 234 of the record for object X to indicate process Pid B-1. Then, as step 4-10, lock manager LMA sends as event E19 a "granted" message to node B, specifically indicating that the lock request (event E15) of process Pid B-1 has been granted.
- 20 Event E20 shows the "granted" message being conveyed by lock manager LMB to process Pid B-1.

- Upon receipt of the "granted message", process Pid B-1 can, at last, enter state 5-9A. In state 5-9A, process
- 25 Pid B-1 sends an "sticky lock request" message to lock managers at all nodes whereat data object X resides in replicated databases. Initially the "sticky lock request" message is sent to local lock manager LMB (event E21), which sends a network message (event E21A) to other nodes. As
- 30 shown associated with event E21 in Fig. 3B, the lingering lock flag has been changed.

-24-

Process Pid B-1 can then obtain access to object X. When process Pid B-1 assumes state 5-3 to release access to data object X (as understood from earlier discussions), an "unlock request" message is sent as event E22 to lock managers at all nodes whereat data object X resides. The lock table supervisor of lock manager LMB clears the object lock flag from field 234 of lock table 220B (step 4-8). After the object lock flag is cleared, lock manager LMB causes its lock schedule queue handler 208 to determine whether any process or node listed in lock schedule queue 222 of node B has requested access to data object X. In the present case, it is noted from queue 222B that process Pid B-2 has made such request, as above mentioned. Accordingly, lock manager LMB causes its lock table supervisor 206 to set the object lock flag in field 234 to indicate that process Pid B-2 now has the lock on data object X (step 4-9). The reference to process Pid B-2 is then removed from the lock scheduler queue 222B of lock manager LMB (step 4-9). Then, at step 4-10 lock manager LMB sends a "granted" message to process Pid B-2 as event E21. The "granted" message of event E22 is responsive to the "lock request" message of event E17 from process Pid B-2.

Although not shown as such in Fig. 3B but understood with respect to the steps of Fig. 4, the events of the preceding paragraph also occur for lock manager LMA of node A. As a result, the lock scheduler queue 222A of node A is also cleared, and the object lock flag in field 232 for object X is set in object lock table 220 to refer to process Pid B-2.

Although for simplicity the invention has been described above in terms of a network 20 including two nodes, i.e., nodes 30A and 30B, it will be readily appreciated that the invention is readily applicable to
5 networks having a greater number of nodes. Moreover, it should be understood that only messages necessary for illustrating the current invention have been illustrated, and that other types of messages including confirmation messages are not illustrated.

10

There is the possibility that dead locks can occur during some scenarios of implementation of the present invention. A dead lock exists when two or more processes are waiting for each other forever. In a dead lock, two
15 processes can end up in waiting queues. Dead lock situations can be remedied by any number of well known techniques, such as the WAIT-DIE algorithm as documented by Rosenkrantz, D.J., Stearns, R.E., and Lewis, P.M., "System Level Concurrency Control for Distributed Database Systems",
20 *ACM Transactions, Database Systems*, 3(2): 178-179, June 1978.

While the invention has been particularly shown and described with reference to the preferred embodiments
25 thereof, it will be understood by those skilled in the art that various alterations in form and detail may be made therein without departing from the spirit and scope of the invention. For example, it should be understood that both the sticky lock flag and the object lock flag for a data
30 object can be maintained in a single table, as opposed to two separate tables as shown in the specifically illustrated embodiments.

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. A node of a computer network, the node comprising:
 - a memory in which a data object is stored, the data object being replicated in other nodes of the network;
 - a lock table;
 - 5 a processor having a lock manager which, when a process native to the node desires access to the data object, sets a lingering lock flag in the lock table for the data object to indicate that the data object can only be accessed at the node to the exclusion of processes of other
 - 10 nodes, the lingering lock flag remaining set after the process native to the node terminates access to the data object, thereby enabling a successive access to the data object by the process native to the node or another process native to the node without notifying other nodes of the
 - 15 network.
2. The apparatus of claim 1, wherein when the process native to the node desires access to the data object, the lock manager sets an object lock flag in the lock table for the data object, and when the process native to the node
- 5 terminates access to the data object, the lock manager clears the object lock flag.
3. The apparatus of claim 1, further comprising a lock scheduling queue for the data object, and wherein when the lock manager receives a request from a non-native process from a second node relating to a desired access of
- 5 the data object at the second node while the lingering lock flag is set for the data object in the lock table of the

node, the lock manager lists the non-native process of the second node as requesting access to the data object in the lock scheduling queue.

4. A computer product comprising a set of programmed instructions stored in a program memory, the set of instructions upon being executed by a processor of a first node of a network, performing the steps of:

5 (1) determining when a process native to the node desires access to a data object, the data object being replicated in other nodes of the network;

 (2) in accordance with step (1), setting a lingering lock flag in a lock table for the data object
10 to indicate that the data object can only be accessed at the node to the exclusion of processes of other nodes; then

 (3) maintaining the lingering lock flag set after the process native to the node terminates access to the data object, thereby enabling a successive access to the data
15 object by the process native to the node or another process native to the node without notifying other nodes of the network.


5. The product of claim 4, further comprising:


 setting an object lock flag in the lock table for the data object when the process native to the node desires access to the data object;

5 clearing the object lock flag when the process native to the node terminates access to the data object.

6. The product of claim 5, further comprising:
maintaining a lock scheduling queue; and
listing a process of the second node as requesting
access to the data object in the lock scheduling queue, when
5 a request is received from the process of the second node
relating to a desired access of the data object at the
second node while the lingering lock flag is set for the
data object in the lock table of the node.
7. A method of regulating access to a data object
replicated across a plurality of nodes of a network, the
method comprising:
when a process native to a first node desires
5 access to the data object, setting a lingering lock flag in
a lock table for the data object to indicate that the data
object can only be accessed at the first node to the
exclusion of processes of other nodes; then
maintaining the lingering lock flag set after the
10 process native to the first node terminates access to the
data object, thereby enabling a successive access to the
data object by the process native to the first node or
another process native to the first node without notifying
other nodes of the network.
8. The method of claim 7, further comprising:
setting an object lock flag in the lock table for
the data object when the process native to the node desires
access to the data object; and
5 clearing the object lock flag when the process
native to the node terminates access to the data object.

9. The method of claim 8, further comprising:
- maintaining a lock scheduling queue; and
 - listing a of the second node as requesting access
- to the data object in the lock scheduling queue, when a
- 5 request is received from the process from the second node
- relating to a desired access of the data object at the
- second node while the lingering lock flag is set for the
- data object in the lock table of the node.

Fig. 1






Fig-3B




Fig. 4

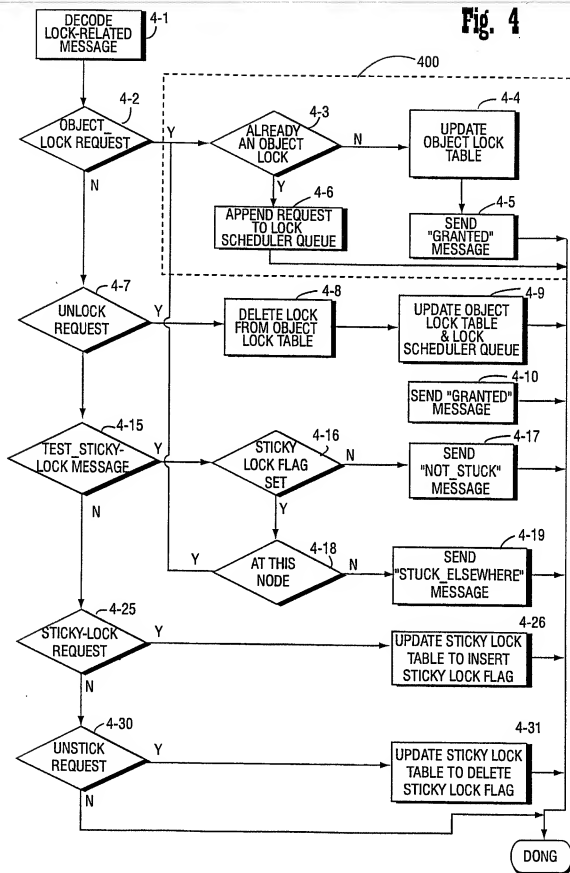


Fig. 5